DOCUMENT RESUME

ED 252 408 SE 045 330

AUTHOR Clancey, William J.
TITLE Classification Problem Solving.
INSTITUTION Stanford Univ., Calif. Dept. of Computer Science.
SPONS AGENCY Office of Naval Research, Arlington, Va. Personnel and Training Research Programs Office.
REPORT NO HPP-84-7; STAN-CS-84-1018
PUB DATE Jul 84
CONTRACT N00014-79C-0302
NOTE 31p.
PUB TYPE Reports - General (140)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Classification; *Computer Science; *Computer Software; *Heuristics; Higher Education; *Problem Solving; Vertical Organization

ABSTRACT

In an attempt to specify in some canonical terms what many heuristic programs known as "expert systems" do, an analysis was made of ten rule-based systems. It was found that these programs proceed through easily identifiable phases of data abstraction, heuristic mapping onto a hierarchy of pre-enumerated solutions, and refinement within this hierarchy: that is, they do classification. Focusing on content rather than representational technology, this paper proposes a set of terms and relations for describing the knowledge used to solve a problem by classification. Then the model used in analyzing the programs MYCIN, SACON, GRUNDY, and SOPHIE III is described and illustrated. These programs have a characteristic inference structure; moreover, this level of description specifies the knowledge needed to solve a problem, independent of its representation in a particular computer language. The classification problem-solving model provides a useful framework for recognizing and representing similar problems, for designing representation tools, and for understanding the problem-solving methods used by non-classification programs. (MNS)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

SE

July 1984

Report No. STAN-CS-84-1018
Also numbered HPP-84-7

ED252408

U.S. DEPARTMENT OF EDUCATION
NATIONAL INSTITUTE OF EDUCATION
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

✓ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official NIE position or policy.

Classification Problem Solving

by

William J. Clancey

Department of Computer Science

Stanford University
Stanford, CA 94305

SE 045 330

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER STAN-CS-84-1018 Tech Report #10	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Classification Problem Solving		5. TYPE OF REPORT & PERIOD COVERED Technical, July 1984
		6. PERFORMING ORG. REPORT NUMBER STAN-CS-84-1018
7. AUTHOR(s) William J. Clancey		8. CONTRACT OR GRANT NUMBER(s) N00014-79C-0302
9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of Computer Science Stanford University Stanford, CA 94305		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR154-482
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, VA 22217		12. REPORT DATE July, 1984 13. NO. OF PAGES 27
14. MONITORING AGENCY NAME & ADDRESS (if diff. from Controlling Office) ONR Representative - Mr. Robin Simpson Durand Aeronautics Building, Rm. 165 Stanford University, Stanford, CA 94305		15. SECURITY CLASS. (of this report) Unclassified 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this report)

Approved for public release: distribution unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from report)

18. SUPPLEMENTARY NOTES

Also HPP Memo 84-7, revision of paper that appeared in the Proceedings of the National Conference on Artificial Intelligence, 1984, pp. 49-55

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

A broad range of heuristic programs--embracing forms of diagnosis, catalog selection, and skeletal planning--accomplish a kind of well-structured problem solving called classification. These programs have a characteristic inference structure that systematically relates data to a preenumerated set of solutions by abstraction, heuristic association, and refinement. This level of description specifies the knowledge needed to solve a problem, independent of its representation in a particular computer language. The classification problem-solving model provides a useful framework for recognizing and representing similar problems, for designing

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

19. KEY WORDS (Continued)

20 ABSTRACT (Continued)

representation tools, and for understanding the problem-solving methods used by non-classification programs.

CLASSIFICATION PROBLEM SOLVING

William J. Clancey

Department of Computer Science
Stanford University, Stanford CA 94305

Contract No. N00014-79C-0302, effective March 15, 1979
Expiration Date: March 14, 1985
Total Amount of Contract -- $1,126,897
Principal Investigator, Bruce G. Buchanan (415) 497-0935
Associate Investigator, William J. Clancey (415) 497-1997

Sponsored jointly by:
Office of Naval Research and Army Research Institute
Personnel and Training Research Programs
Psychological Sciences Division
Contract Authority No. NR 154-482
Scientific Officer: Dr. Marshall Farr

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either express or implied, or the Office of Naval Research or the U.S. Government.

Approved for public release; distribution unlimited. Reproduction in whole or in part is permitted for any purpose of the United States Government.

Abstract

A broad range of heuristic programs—embracing forms of diagnosis, catalog selection, and skeletal planning—accomplish a kind of well-structured problem solving called classification. These programs have a characteristic inference structure that systematically relates data to a pre-enumerated set of solutions by abstraction, heuristic association, and refinement. This level of description specifies the knowledge needed to solve a problem, independent of its representation in a particular computer language. The classification problem-solving model provides a useful framework for recognizing and representing similar problems, for designing representation tools, and for understanding the problem-solving methods used by non-classification programs.

1. Introduction

Over the past decade a variety of heuristic programs have been written to solve problems in diverse areas of science, engineering, business, and medicine. Yet, presented with a given "knowledge engineering tool," such as EMYCIN [39], we are still hard-pressed to say what kinds of problems it can be used to solve well. Various studies have demonstrated advantages of using one representation language instead of another—for ease in specifying knowledge relationships, control of reasoning, and perspicuity for maintenance and explanation [9, 38, 1, 2, 10]. Other studies have characterized in low-level terms why a given problem might be inappropriate for a given language, for example, because data are time-varying or subproblems interact [21]. But attempts to describe *kinds of problems* in terms of shared features have not been entirely satisfactory: Applications-oriented descriptions like "diagnosis" are too general (e.g., the program might not use a device model), and technological terms like "rule-based" don't describe what problem is being solved [18, 19]. Logic has been suggested as a tool for a "knowledge level" analysis that would specify what a heuristic program does, independent of its implementation in a programming language [30, 29]. However, we have lacked a set of terms and relations for doing this.

In an attempt to specify in some canonical terms what many heuristic programs known as "expert systems" do, an analysis was made of ten rule-based systems (including MYCIN, SACON, and The Drilling Advisor), a frame-based system (GRUNDY) and a program coded directly in LISP (SOPHIE III). There is a striking pattern: These programs proceed through easily identifiable phases of data abstraction, heuristic mapping onto a hierarchy of pre-enumerated solutions, and refinement within this hierarchy. In short, these programs do what is commonly called *classification*.

Focusing on content rather than representational technology, this paper proposes a set of terms and relations for describing the knowledge used to solve a problem by classification. Subsequent

sections describe and illustrate the classification model in the analysis of MYCIN, SACON, GRUNDY, and SOPHIE III. Significantly, a knowledge level description of these programs corresponds very well to psychological models of expert problem solving. This suggests that the classification problem solving model captures general principles of how experiential knowledge is organized and used, and thus generalizes some cognitive science results. There are several strong implications for the practice of building expert systems and continued research in this field.

2. Classification problem solving defined

We develop the idea of classification problem solving by starting with the common sense notion and relating it to the reasoning that occurs in heuristic programs.

2.1. Simple classification

As the name suggests, the simplest kind of classification problem is to identify some unknown object or phenomenon as a member of a known class of objects or phenomena. Typically, these classes are stereotypes that are hierarchically organized, and the process of identification is one of matching observations of an unknown entity against features of known classes. A paradigmatic example is identification of a plant or animal, using a guidebook of features, such as coloration, structure, and size.

Some terminology we will find helpful: The *problem* is the object or phenomenon to be identified; *data* are observations describing this problem; possible *solutions* are patterns (variously called schema, frames or units); each solution has a set of *features* (slots or facets) that in some sense describe the concept either categorically or probabilistically; solutions are grouped into *a specialization hierarchy* based on their features (in general, not a single hierarchy, but multiple, directed acyclic graphs); a *hypothesis* is a solution that is under consideration; *evidence* is data that partially matches some hypothesis; the *output* is some solution.

The essential characteristic of a classification problem is that the problem solver selects from a set of pre-enumerated solutions. This does not mean, of course, that the "right answer" is necessarily one of these solutions, just that the problem solver will only attempt to match the data against the known solutions, rather than construct a new one. Evidence can be uncertain and matches partial, so the output might be a ranked list of hypotheses.

Besides matching, there are several *rules of inference* for making assertions about solutions. For example, evidence for a class is indirect evidence that one of its subtypes is present. Conversely,

given a closed world assumption, evidence against all of the subtypes is evidence against a class. *Search operators* for finding a solution also capitalize on the hierarchical structure of the solution space. These operators include: *refining* a hypothesis to a more specific classification; *categorizing* the problem by considering superclasses of partially matched hypotheses; and *discriminating* among hypotheses by contrasting their superclasses [31, 32, 12]. For simplicity, we will refer to the entire process of applying these rules of inference and operators as *refinement*. The specification of this process—a control strategy—is an orthogonal issue which we will consider later.

2.2. Data abstraction

In the simplest problems, data are solution features, so the matching and refining process is direct. For example, an unknown organism in MYCIN can be classified directly given the supplied data of gram stain and morphology.

For many problems, solution features are not supplied as data, but are inferred by *data abstraction*. There are three basic relations for abstracting data in heuristic programs:

- *qualitative abstraction* of quantitative data ("if the patient is an adult and white blood count is less than 2500, then the white blood count is low");
- *definitional abstraction* ("if the structure is one-dimensional of network construction, then its shape is a beam"); and
- *generalization* in a subtype hierarchy ("if the client is a judge, then he is an educated person").

These interpretations are usually made by the program with certainty; thresholds and qualifying contexts are chosen so the conclusion is categorical. It is common to refer to this knowledge as "descriptive," "factual," or "definitional."

2.3. Heuristic classification

In simple classification, the data may directly match the solution features or may match after being abstracted. In heuristic classification, solution features may also be matched heuristically. For example, MYCIN does more than identify an unknown organism in terms of features of a laboratory culture: It heuristically relates an abstract characterization of the patient to a classification of diseases. We show this *inference structure* schematically, followed by an example (Figure 2-1).

Basic observations about the patient are abstracted to patient categories, which are heuristically linked to diseases and disease categories. While only a subtype link with E.coli infection is shown

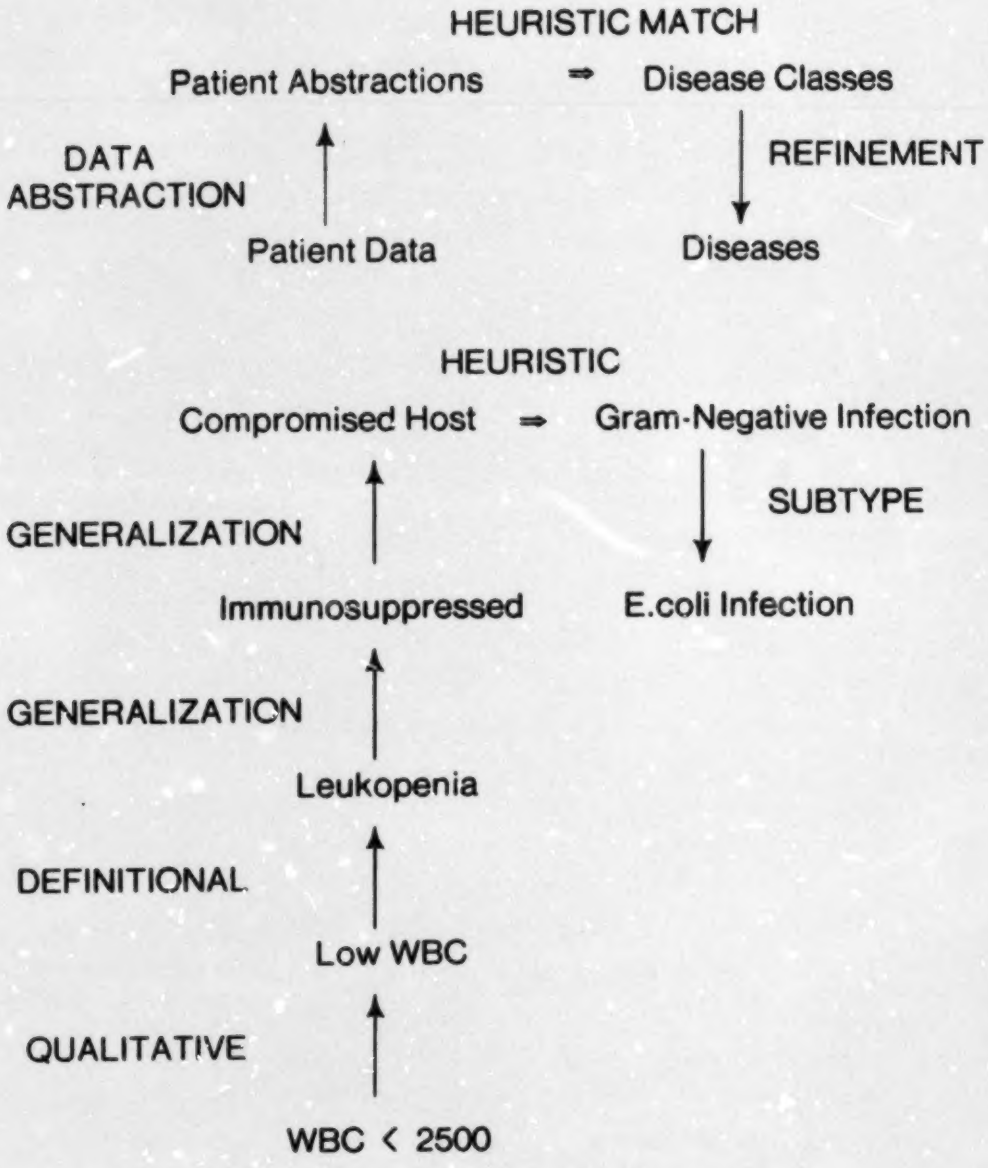

Figure 2-1: Inference structure of MYCIN

here, evidence may actually derive from a combination of inferences. Some data might directly match E.coli by identification Discrimination with competing subtypes of gram-negative infection might also provide evidence. As stated earlier, the order in which these inferences are made is a matter of control strategy.

The important link we have added is a heuristic association between a characterization of the patient ("compromised host") and categories of diseases ("gram-negative infection"). Unlike the factual and hierarchical evidence propagation we have considered to this point, this inference makes a great leap. A *heuristic relation* is based on some implicit, possibly incomplete, model of the world. This relation is often empirical, based just on experience; it corresponds most closely to the "rules of thumb" often associated with heuristic programs [14].

Heuristics of this type reduce search by skipping over intermediate relations (this is why we don't call abstraction relations "heuristics"). These associations are usually uncertain because the intermediate relations may not hold in the specific case. Intermediate relations may be omitted because they are unobservable or poorly understood. In a medical diagnosis program, heuristics typically skip over the causal relations between symptoms and diseases.

To repeat, classification problem solving involves heuristic association of an abstracted problem statement onto features that characterize a solution. This can be shown schematically in simple terms (Figure 2-2).

This diagram summarizes how a distinguished set of terms (data, data abstractions, solution abstractions, and solutions) are related systematically by *different* kinds of relations and rules of inference. This is the *structure of inference* in classification problem solving.

In a study of physics problem solving, Chi [8] calls data abstractions "transformed" or "second order problem features." In an important and apparently common variant of the simple model, data abstractions are themselves patterns that are heuristically matched. In essence, there is a sequence of classification problems. GRUNDY, analyzed below, illustrates this.

2.4. Search in classification problem solving

The issue of search is orthogonal to the kinds of inference we have been considering. "Search" refers to how a network made up of abstraction, heuristic, and refinement relations is interpreted, how the flow of inference actually might proceed in solving a problem. Following Hayes [18], we call this the *process structure*. There are three basic process structures in classification problem solving:

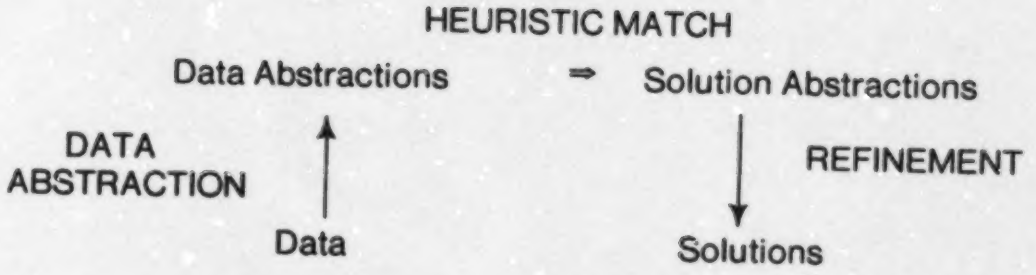

Figure 2-2: Classification problem solving inference structure

1. *Data-directed search*: The program works forwards from data to abstractions, matching solutions until all possible (or non-redundant) inferences have been made.

2. *Solution- or Hypothesis-directed search*: The program works backwards from solutions, collecting evidence to support them, working backwards through the heuristic relations to the data abstractions and required data to solve the problem. If solutions are hierarchically organized, then categories are considered before direct features of more specific solutions.

3. *Opportunistic search*: The program combines data and hypothesis-directed reasoning [20]. Data abstraction rules tend to be applied immediately as data become available. Heuristic rules "trigger" hypotheses, followed by a focused, hypothesis-directed search. New data may cause refocusing. By reasoning about solution classes, search need not be exhaustive.

Data- and hypothesis-directed search are not to be confused with the implementation terms "forward" or "backward chaining." R1 provides a superb example of how different implementation and knowledge level descriptions can be. Its rules are *interpreted* by forward-chaining, but it does a

form of hypothesis-directed search, systematically setting up subproblems by a fixed procedure that focuses reasoning on spatial subcomponents of a solution [28].

The degree to which search is *focused* depends on the level of indexing in the implementation and how it is exploited. For example, MYCIN's "goals" are solution classes (e.g., types of bacterial meningitis), but selection of rules for specific solutions (e.g., E.coli meningitis) is unordered. Thus, MYCIN's search within each class is unfocused [11].

The choice of process structure depends on the number of solutions, whether they can be categorically constrained, usefulness of data (the density of rows in a data/solution matrix), and the cost for acquiring data.

3. Examples of classification problem solving

Here we schematically describe the architectures of SACON, GRUNDY, and SOPHIE III in terms of classification problem solving. These are necessarily very brief descriptions, but reveal the value of this kind of analysis by helping us to understand what the programs do. After a statement of the problem, the general inference structure and an example inference path are given, followed by a brief discussion.

3.1. SACON

Problem: SACON [3] selects classes of behavior that should be further investigated by a structural analysis simulation program (Figure 3-1).

Discussion: SACON solves two problems by classification—analyzing a structure and then selecting a program. It begins by heuristically selecting a simple numeric model for analyzing a structure (such as an airplane wing). The model produces stress and deflection estimates, which the program then abstracts in terms of features that hierarchically describe different configurations of the MARC simulation program. There is no refinement because the solutions to the first problem are just a simple set of possible models, and the second problem is only solved to the point of specifying program classes. (In another software configuration system we analyzed, specific program input parameters are inferred in a refinement step.)

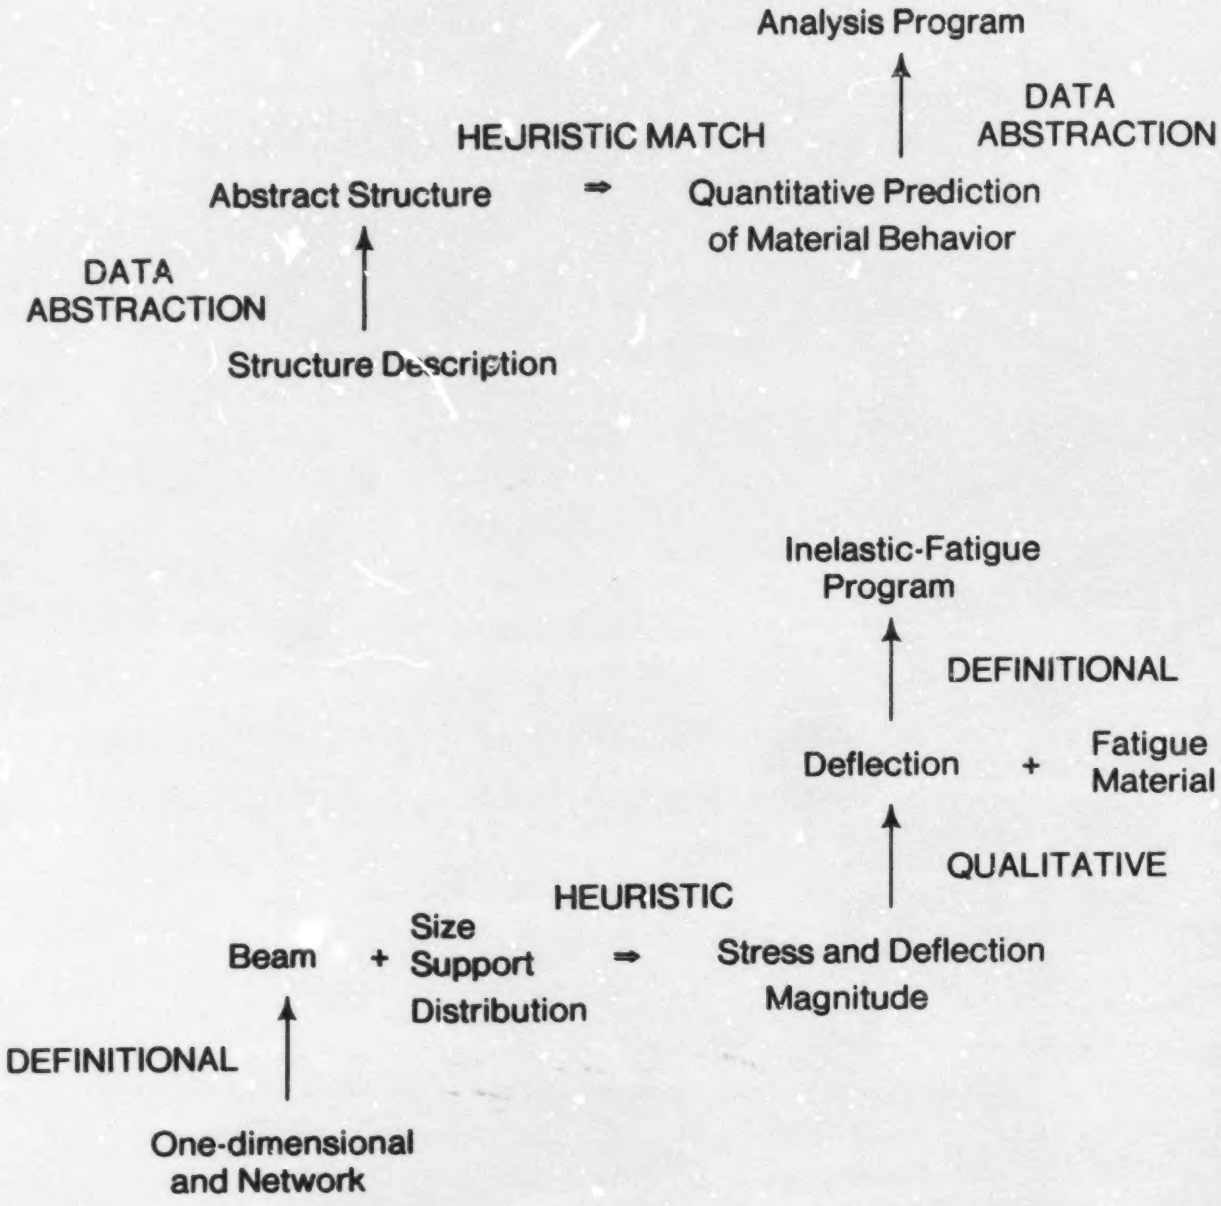

Figure 3-1: Inference structure of SACON

3.2. GRUNDY

Problem: GRUNDY [33] heuristically classifies a reader's personality and selects books he might like to read (Figure 3-2).

Discussion: GRUNDY solves two classification problems heuristically. Illustrating the power of a knowledge level analysis, we discover that the people and book classifications are not distinct in the implementation. For example, "fast plots" is a book characteristic, but in the implementation "likes fast plots" is associated with a person stereotype. The relation between a person stereotype and "fast plots" is heuristic and should be distinguished from abstractions of people and books. One objective of the program is to learn better people stereotypes (user models). The classification description of the user modeling problem shows that GRUNDY should also be learning better ways to characterize books, as well as improving its heuristics. If these are not treated separately, learning may be hindered. This example illustrates why a knowledge level analysis should precede representation.

It is interesting to note that GRUNDY does not attempt to perfect the user model before recommending a book. Rather, refinement of the person stereotype occurs when the reader rejects book suggestions. Analysis of other programs indicates that this multiple-pass process structure is common. For example, the Drilling Advisor makes two passes on the causes of sticking, considering general, inexpensive data first, just as medical programs commonly consider the "history and physical" before laboratory data.

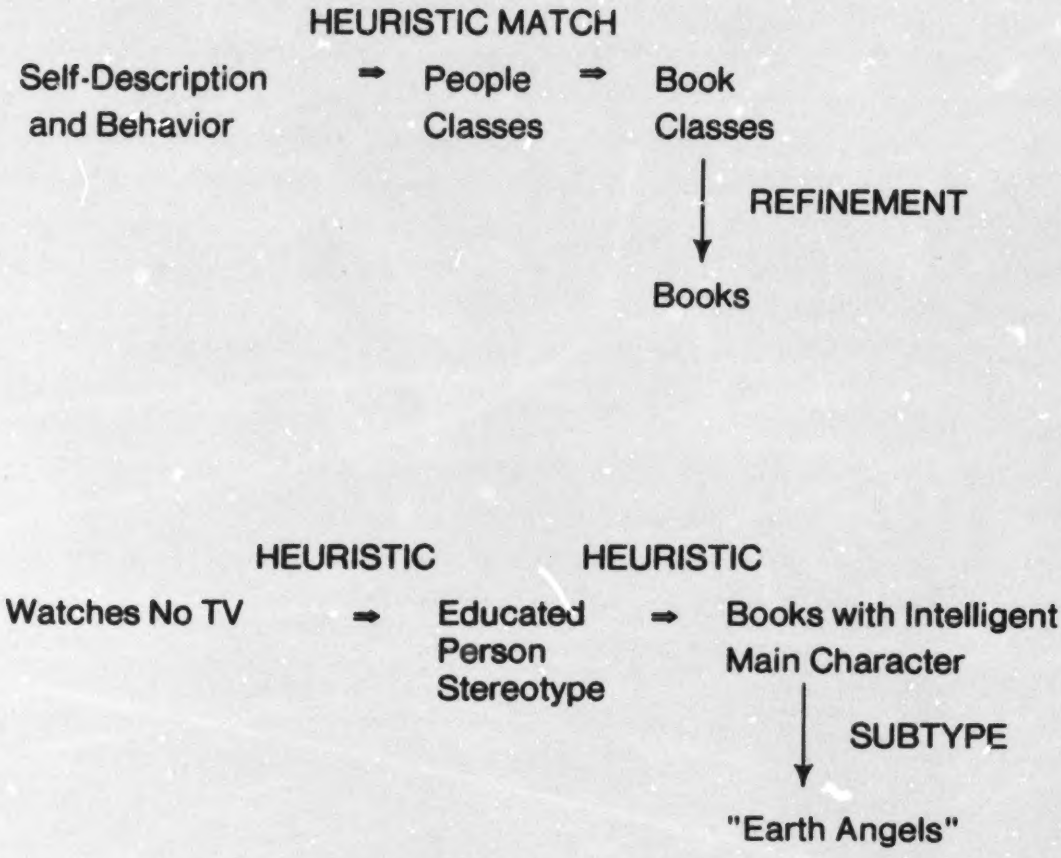

Figure 3-2: Inference structure of GRUNDY

3.3. SOPHIE III

Problem: SOPHIE III [5] classifies an electronic circuit in terms of the component that is causing faulty behavior (Figure 3-3).

Discussion: SOPHIE's set of pre-enumerated solutions is a lattice of valid and faulty circuit behaviors. In contrast with MYCIN, solutions are device states and component flaws, not stereotypes of disorders, and they are related causally, not by features. Data are not just external device behaviors, but include internal component measurements propagated by the causal analysis of the LOCAL program. Reasoning about assumptions plays a central role in matching hypotheses. In spite of these differences, the inference structure of abstractions, heuristic relations, and refinement fits the classification model, demonstrating its generality and usefulness for describing complex reasoning.

4. Causal process classification

To further illustrate the value of a knowledge level analysis, we describe a generic inference structure common to medical diagnosis programs, which we call *causal process classification*, and use it to contrast the goals of electronic circuit and medical diagnosis programs.

In SOPHIE, valid and abnormal device states are exhaustively enumerated, can be directly confirmed, and are causally related to component failures. None of this is generally possible in medical diagnosis, nor is diagnosis in terms of component failures alone sufficient for selecting therapy. Medical programs that deal with multiple disease processes (unlike MYCIN) do reason about abnormal states (called *pathophysiologic states*, e.g., "increased pressure in the brain"), directly analogous to the abnormal states in SOPHIE. But curing an illness generally involves determining the cause of the component failure. These "final causes" (called diseases, syndromes, etiologies) are processes that affect the normal functioning of the body (e.g., trauma, infection, toxic exposure, psychological disorder). Thus, medical diagnosis more closely resembles the task of computer system diagnosis in considering how the body relates to its environment [25]. In short, there are two problems: First to explain symptoms in terms of abnormal internal states, and second to explain this behavior in terms of external influences (as well as congenital and degenerative component flaws). This is the inference structure of programs like CASNET [42] and NEOMYCIN [9] (Figure 4-1).

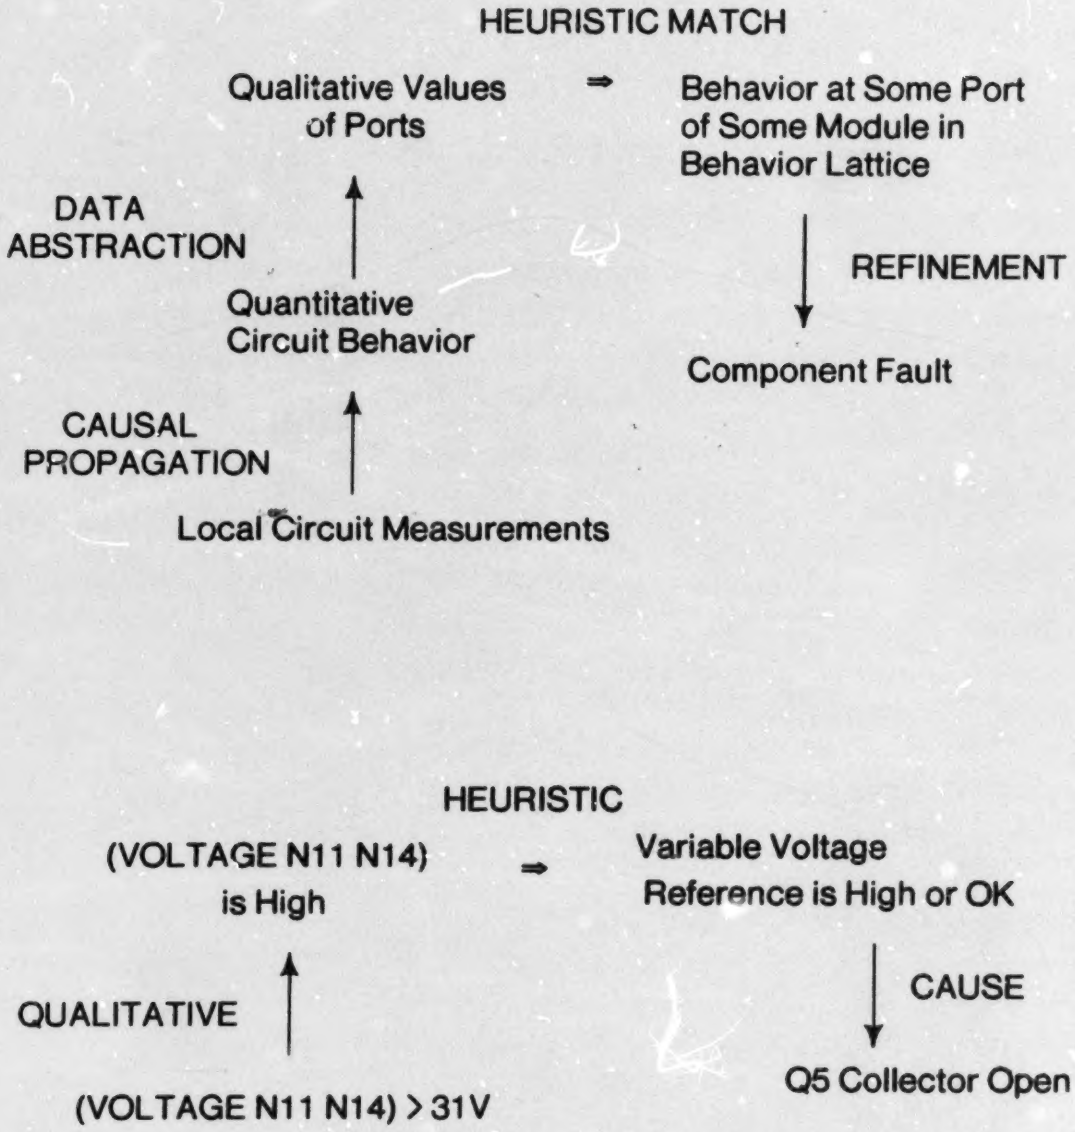

Figure 3-3: Inference structure of SOPHIE

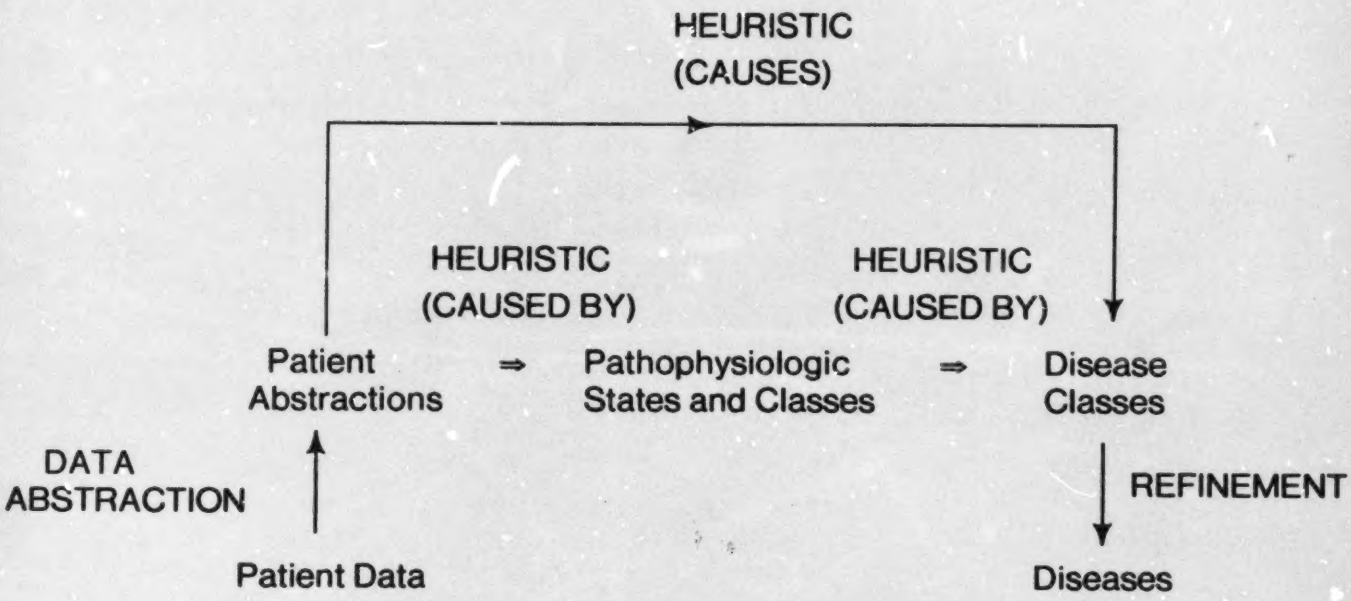

Figure 4-1: Inference structure of causal process classification

A network of causally related pathophysiologic states causally relates data to diseases[1]. The causal relations are themselves heuristic because they assume certain physiologic structure and behavior, which is often poorly understood and not represented. In contrast with pathophysiologic states, diseases are abstractions of *processes*--causal stories with agents, locations, and sequences of events. Disease networks are organized by these process features (e.g., an organ system taxonomy organizes diseases by location). A more general term for disease is *disorder stereotype*. In *process control* problems, such as chemical manufacturing, the most general disorder stereotypes correspond to stages in a process (e.g., mixing, chemical reaction, filtering, packaging). Subtypes

[1]Programs differ in whether they treat pathophysiologic states as independent solutions (NEOMYCIN) or find the causal path that best accounts for the data (CASNET). Moreover, a causal explanation of the data requires finding a state network, including normal states, that is internally consistent on multiple levels of detail. Combinatorial problems, as well as elegance, argue against pre-enumerating solutions, so such a network must be constructed, as in ABEL [31]. In SOPHIE, the LOCAL program deals with most of the state interactions at the component level, others are captured in the exhaustive hierarchy of module behaviors. A more general solution is to use a structure/function device model and general diagnostic operators, as in DART [16].

correspond to what can go wrong at each stage [12].

To summarize, a knowledge level analysis reveals that medical and electronic diagnosis programs are not all trying to solve the same kind of problem. Examining the nature of solutions, we see that in a electronic circuit diagnosis program like SOPHIE solutions are component flaws. Medical diagnosis programs like CASNET attempt a second step, *causal process classification*, which is to explain abnormal states and flaws in terms of processes external to the device or developmental processes affecting its structure. It is this experiential knowledge—what can affect the device in the world—that is captured in disease stereotypes. This knowledge can't simply be replaced by a model of device structure and function, which is concerned with a different level of analysis.

5. What is non-classification problem solving?

We first summarize the applications we have considered by observing that all classification problem solving involves *selection* of a solution. We can characterize kinds of problems by what is being selected:

- *diagnosis*: solutions are faulty components (SOPHIE) or processes affecting the device (MYCIN);
- *user model*: solutions are people stereotypes in terms of their goals and beliefs (first phase of GRUNDY);
- *catalog selection*: solutions are products, services, or activities, e.g., books, personal computers, careers, travel tours, wines, investments (second phase of GRUNDY);
- *theoretical analysis*: solutions are numeric models (first phase of SACON);
- *skeletal planning*: solutions are plans, such as packaged sequences of programs and parameters for running them (second phase of SACON, also first phase of experiment planning in MOLGEN [15]).

A common misconception is that the description "classification problem" is an inherent property of a problem, opposing, for example, classification with design [37]. However, classification problem solving, as defined here, is *a description of how a problem is solved by a particular problem solver.* If the problem solver has a priori knowledge of solutions and can relate them to the problem description by data abstraction, heuristic association, and refinement, then the problem can be solved by classification. For example, if it were practical to enumerate all of the computer configurations R1 might select, or if the solutions were restricted to a predetermined set of designs, the program could be reconfigured to solve its problem by classification.

Furthermore, as illustrated by ABEL, it is incorrect to say that medical diagnosis is a "classification problem." Only *routine* medical diagnosis problems can be solved by classification [32]. When there are multiple, interacting diseases, there are too many possible combinations for the problem solver to have considered them all before. Just as ABEL reasons about interacting states, the physician must construct a consistent network of interacting diseases to explain the symptoms. The problem solver *formulates a solution*; he doesn't just make yes-no decisions from a set of fixed alternatives. For this reason, Pople calls non-routine medical diagnosis an ill-structured problem [36] (though it may be more appropriate to reserve this term for the theory formation task of the physician-scientist who is defining new diseases).

In summary, a useful distinction is whether a solution is selected or constructed. To *select* a solution, the problem solver needs experiential ("expert") knowledge in the form of *patterns of problems and solutions* and heuristics relating them. To *construct* a solution, the problem solver applies models of structure and behavior, by which objects can be assembled, diagnosed, or employed in some plan.

Whether the solution is taken off the shelf or is pieced together has important computational implications for choosing a representation. In particular, construction problem-solving methods such as constraint propagation and dependency-directed backtracking have data structure requirements that may not be easily satisfied by a given representation language. For example—returning to a question posed in the introduction—applications of EMYCIN are generally restricted to problems that can be solved by classification.

6. Knowledge level analysis

As a set of terms and relations for describing knowledge (e.g, data, solutions, kinds of abstraction, refinement operators, the meaning of "heuristic"), the classification model provides a *knowledge level analysis* of programs, as defined by Newell [29]. It "serves as a specification of what a reasoning system should be able to do." Like a specification of a conventional program, this description is distinct from the representational technology used to implement the reasoning system. Newell cites Schank's conceptual dependency structure as an example of a knowledge level analysis. It indicates "what knowledge is required to solve a problem... how to encode knowledge of the world in a representation."

After a decade of "explicitly" representing knowledge in AI languages, it is ironic that the pattern of classification problems should have been so difficult to see. In retrospect, certain views were

emphasized at the expense of others:

- *Procedureless languages.* In an attempt to distinguish heuristic programming from traditional programming, procedural constructs are left out of representation languages (such as EMYCIN, OPS, KRL [26]). Thus, inference relations cannot be stated separately from how they are to be used [18, 19].

- *Heuristic nature of problem solving.* Heuristic association has been emphasized at the expense of the relations used in data abstraction and refinement. In fact, some expert systems do only simple classification: they have no heuristics or "rules of thumb," the key idea that is supposed distinguish this class of computer programs.

- *Implementation terminology.* In emphasizing new implementation technology, terms such as "modular" and "goal directed" were more important to highlight than the content of the programs. In fact, "goal directed" characterizes any rational system and says very little about how knowledge is used to solve a problem. "Modularity" is a representational issue of indexing.

Nilsson has proposed that logic should be the *lingua franca* for knowledge level analysis [30]. Our experience with the classification model suggests that the value of using logic is in adopting a set of terms and relations for describing knowledge (e.g., kinds of abstraction). Logic is valuable as a tool for knowledge level analysis because it emphasizes *relations*, not just implication.

While rule-based languages do not make important knowledge level distinctions, they have nevertheless provided an extremely successful programming framework for classification problem solving. Working backwards (backchaining) from a pre-enumerated set of solutions guarantees that only the relevant rules are tried and useful data considered. Moreover, the program designer is encouraged to use means-ends analysis, a clear framework for organizing rule writing.

7. Related analyses

Several researchers have described portions of the classification problem solving model, influencing this analysis. For example, in CRYSALIS [13] data and hypothesis abstraction are clearly separated. The EXPERT rule language [40] similarly distinguishes between "findings" and a taxonomy of hypotheses. In PROSPECTOR [17], rules are expressed in terms of relations in a semantic network. In CENTAUR [2], a variant of MYCIN, solutions are explicitly *prototypes* of diseases. Chandrasekaran and his associates have been strong proponents of the classification model: "The normal problem-solving activity of the physician... (is) a process of classifying the case as an element of a disease taxonomy" [7]. Recently, Chandrasekaran and Weiss and Kulikowski have generalized the classification schemes used by their programs (MDX and EXPERT) to characterize problems solved by other expert systems [6, 41].

A series of knowledge representation languages beginning with KRL have identified structured abstraction and matching as a central part of problem solving [4]. Building on the idea that "frames" are not just a computational construct, but a theory about a kind of knowledge [19], cognitive science studies have described problem solving in terms of classification. For example, routine physics problem solving is described by Chi [8] as a process of data abstraction and heuristic mapping onto solution schemas ("experts cite the abstracted features as the relevant cues (of physics principles)"). The inference structure of SACON, heuristically relating structural abstractions to numeric models, is the same.

Related to the physics problem solving analysis is a very large body of research on the nature of schemas and their role in understanding [35, 34]. More generally, the study of classification, particularly of objects, also called *categorization*, has been a basic topic in psychology for several decades (e.g., see the chapter on "conceptual thinking" in [22]). However, in psychology the emphasis has been on the nature of categories and how they are formed (an issue of learning). The programs we have considered make an identification or selection from a pre-existing classification (an issue of memory retrieval). In recent work, Kolodner combines the retrieval and learning process in an expert system that learns from experience [23]. Her program uses the MOPS representation, a classification model of memory that interleaves generalizations with specific facts [24].

8. Conclusions

A wide variety of problems can be described in terms of heuristic mapping of data abstractions onto a fixed, hierarchical network of solutions. This problem solving model is supported by psychological studies of human memory and the role of classification in understanding. There are significant implications for expert systems research:

- The model provides a *high-level structure for decomposing problems*, making it easier to recognize and represent similar problems. For example, problems can be characterized in terms of sequences of classification problems. Catalog selection programs might be improved by incorporating a more distinct phase of user modelling, in which needs or requirements are classified first. Diagnosis programs might profitably make a stronger separation between device-history stereotypes and disorder knowledge. A generic knowledge engineering tool can be designed specifically for classification problem solving. The advantages for knowledge acquisition carry over into explanation and teaching.

- The model provides a *basis for choosing application problems*. For example, problems can be selected that will teach us more about the nature of abstraction and how other forms of inference (e.g., analogy, simulation, constraint posting) are combined with classification.

- The model provides a *foundation for describing representation languages* in terms of epistemologic adequacy [27], so that the leverage they provide can be better understood. For example, for classification it is advantageous for a language to provide constructs for representing problem solutions as a network of schemas.

- The model provides a *focus for cognitive studies* of human categorization of knowledge and search strategies for retrieval and matching, suggesting principles that might be used in expert programs. Learning research might similarly focus on the inference and process structure of classification problem solving.

Finally, it is important to remember that expert systems are programs. Basic computational ideas such as input, output, and sequence, are essential for describing what they do. The basic methodology of our study has been to ask, "What does the program conclude about? How does it get there from its input?" We characterize the flow of inference, identifying data abstractions, heuristics, implicit models and assumptions, and solution categories along the way. If heuristic programming is to be different from traditional programming, a knowledge level analysis should always be pursued to the deepest levels of our understanding, even if practical constraints prevent making explicit in the implemented program everything that we know. In this way, knowledge *engineering* can be based on sound principles that unite it with studies of cognition and representation.

Acknowledgments

Many of these ideas were stimulated by discussions with Denny Brown in our attempt to develop a framework for teaching knowledge engineering. I would also like to thank Tom Dietterich, Steve Hardy, and Peter Szolovits for their suggestions and encouragement.

The Drilling Advisor is a product of Teknowledge, Inc.

This research has been supported in part by ONR and ARI Contract N00014-79C-0302. Computational resources have been provided by the SUMEX-AIM facility (NIH grant RR00785).

References

[1] Aiello, N.
A comparative study of control strategies for expert systems: AGE implementation of three variations of PUFF.
In *Proceedings of the National Conference on AI*, pages 1-4. Washington, D.C., August, 1983.

[2] Aikins J. S.
Prototypical knowledge for expert systems.
Artificial Intelligence 20(2):163-210, 1983.

[3] Bennett, J., Creary, L., Englemore, R., and Melosh, R.
SACON: A knowledge-based consultant for structural analysis.
STAN-CS-78-699 and HPP Memo 78-23, Stanford University, Sept, 1978.

[4] Bobrow, D. and Winograd, T.
KRL: Another perspective.
Cognitive Science 3:29-42, 1979.

[5] Brown, J. S., Burton, R. R., and de Kleer, J.
Pedagogical, natural language, and knowledge engineering techniques in SOPHIE I, II, and III.
In D. Sleeman and J. S. Brown (editors), *Intelligent Tutoring Systems*, pages 227-282. Academic Press, 1982.

[6] Chandrasekaran, B.
Expert systems: Matching techniques to tasks.
In W. Reitman (editor), *AI Applications for Business*, pages 116-132. Ablex Publishing Corp., 1984.

[7] Chandrasekaran, B. and Mittal, S.
Conceptual representation of medical knowledge.
In M. Yovits (editor), *Advances in Computers*, pages 217-293. Academic Press, New York, 1983.

[8] Chi, M. T. H., Feltovich, P. J., Glaser, R.
Categorization and representation of physics problems by experts and novices.
Cognitive Science 5:121-152, 1981.

[9] Clancey, W. J. and Letsinger, R.
NEOMYCIN: Reconfiguring a rule-based expert system for application to teaching.
In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 829-836. August, 1981.
(Revised version to appear in Clancey and Shortliffe (editors), *Readings in Medical Artificial Intelligence: The First Decade*, Addison-Wesley, 1983).

[10] Clancey, W. J.
The advantages of abstract control knowledge in expert system design.
In *Proceedings of the National Conference on AI*, pages 74-78. Washington, D.C., August, 1983.

[11] Clancey, W. J.
The epistemology of a rule-based expert system: A framework for explanation.
Artificial Intelligence 20(3):215-251, 1983.

[12] Clancey, W. J.
Acquiring, representing, and evaluating a competence model of diagnosis.
HPP Memo 84-2, Stanford University, February, 1984.
(To appear in Chi, Glaser, and Farr (Eds.), *The Nature of Expertise*, in preparation.).

[13] Engelmore, R. and Terry, A.
Structure and function of the CRYSALIS system.
In *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pages 250-256. August, 1979.

[14] Feigenbaum, E. A.
The art of artificial intelligence: I. Themes and case studies of knowledge engineering.
In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 1014-1029. August, 1977.

[15] Friedland, P. E.
Knowledge-based experiment design in molecular genetics.
Technical Report STAN-CS-79-771, Stanford University, October, 1979.

[16] Genesereth, M. R.
Diagnosis using hierarchical design models.
In *Proceedings of the National Conference on AI*, pages 278-283. Pittsburgh, PA, August, 1982.

[17] Hart, P. E.
Observations on the development of expert knowledge-based systems.
In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 1001-1003. August, 1977.

[18] Hayes, P.J.
In defence of logic.
In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 559-565. August, 1977.

[19] Hayes, P.
The logic of frames.
In D. Metzing (editor), *Frame Conceptions and Text Understanding*, pages 45-61. de Gruyter, 1979.

[20] Hayes-Roth, B. and Hayes-Roth, F.
A cognitive model of planning.
Cognitive Science 3:275-310. 1979.

[21] Hayes-Roth, F., Waterman, D., and Lenat, D. (eds.).
Building expert systems.
Addison-Wesley, New York, 1983.

[22] Johnson-Laird, P. N. and Wason, P. C.
Thinking: Readings in Cognitive Science.
Cambridge University Press, Cambridge, 1977.

[23] Kolodner, J. L.
The role of experience in development of expertise.
In *Proceedings of the National Conference on AI*, pages 273-277. Pittsburgh, PA, August, 1982.

[24] Kolodner, J.
Maintaining organization in a dynamic long-term memory.
Cognitive Science 7:243-280, 1983.

[25] Lane, W. G.
Input/output processing.
In Stone, H. S. (editor), *Introduction to Computer Architecture, 2nd Edition*, chapter 6, . Science Research Associates, Inc., Chicago, 1980.

[26] Lehnert, W., and Wilks, Y.
A critical perspective on KRL.
Cognitive Science 3:1-28, 1979.

[27] McCarthy, J. and Hayes, P.
Some philosophical problems from the standpoint of Artificial Intelligence.
In B. Meltzer and D. Michie (editors), *Machine Intelligence 4*, pages 463-502. Edinburgh University Press, 1969.

[28] McDermott, J.
R1: A rule-based configurer of computer systems.
Artificial Intelligence 19(1):39-88, 1982.

[29] Newell, A.
The knowledge level.
Artificial Intelligence 18(1):87-127, 1982.

[30] Nilsson, N. J.
The interplay between theoretical and experimental methods in Artificial Intelligence.
Cognition and Brain Theory 4(1):69-74, 1981.

[31] Patil, R. S., Szolovits, P., and Schwartz, W. B.
Causal understanding of patient illness in medical diagnosis.
In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 893-899. August, 1981.

[32] Pople, H.
Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnostics.
In P. Szolovits (editor), *Artificial Intelligence in Medicine*, pages 119-190. Westview Press, 1982.

[33] Rich, E.
User modeling via stereotypes.
Cognitive Science 3:355-366, 1979.

[34] Rumelhart, D. E. and Norman, D. A.
Representation in memory.
Technical Report CHIP-116, Center for Human Information Processing, University of California, June, 1983.

[35] Schank, R. C., and Abelson, R. P.
Scripts, Plans, Goals, and Understanding.
Lawrence Erlbaum Associates, Hillsdale, NJ, 1975.

[36] Simon, H. A.
The structure of ill structured problems.
Artificial Intelligence 4:181-201, 1973.

[37] Sowa, J. F.
Conceptual Structures.
Addison-Wesley, Reading, MA, 1984.

[38] Swartout W. R.
Explaining and justifying in expert consulting programs.
In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 815-823. August, 1981.

[39] van Melle, W.
A domain-independent production rule system for consultation programs.
In *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pages 923-925. August, 1979.

[40] Weiss, S. M. and Kulikowski, C. A.
EXPERT: A system for developing consultation models.
In *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pages 942-947. August, 1979.

[41] Weiss, S. M. and Kulikowski, C. A.
A Practical Guide to Designing Expert Systems.
Rowman and Allanheld, Totowa, NJ, 1984.

[42] Weiss, S. M., Kulikowski, C. A., Amarel, S., and Safir, A.
A model-based method for computer-aided medical decision making.
Artificial Intelligence 11:145-172, 1978.

1 Dr. Genevieve Haddad
Program Manager
Life Sciences Directorate
AFOSR
Bolling AFB, DC 20332

1 Dr. John Tangney
AFOSR/NL
Bolling AFB, DC 20332

1 Dr. Joseph Yasatuke
AFHRL/LRT
Lowry AFB, CO 80230

Department of Defense

12 Defense Technical Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC

1 Military Assistant for Training and
Personnel Technology,
Office of the Under Secretary of Defense
for Research & Engineering
Room 3D129, The Pentagon
Washington, DC 20301

1 Major Jack Thorpe
DARPA
1400 Wilson Blvd.
Arlington, VA 22209

Civilian Agencies

1 Dr. Patricia A. Butler
NIE-BRN Bldg, Stop # 7
1200 19th St., NW
Washington, DC 20208

1 Dr. Susan Chipman
Learning and Development
National Institute of Education
1200 19th Street NW
Washington, DC 20208

1 Edward Esty
Department of Education, OERI
MS 40
1200 19th St., NW
Washington, DC 20208

1 Dr. Arthur Melmed
724 Brown
U. S. Dept. of Education
Washington, DC 20208

1 Dr. Andrew R. Molnar
Office of Scientific and Engineering
Personnel and Education
National Science Foundation
Washington, DC 20550

1 Dr. Everett Palmer
Mail Stop 239-3
NASA-Ames Research Center
Moffett Field, CA 94035

1 Dr. Mary Stoddard
C 10, Mail Stop B296
Los Alamos National Laboratories
Los Alamos, NM 87545

1 Chief, Psychological Research Branch
U. S. Coast Guard (G-P-1/2/TP42)
Washington, DC 20593

1 Dr. Edward C. Weiss
National Science Foundation
1800 G Street, NW
Washington, DC 20550

1 Dr. Frank Withrow
U. S. Office of Education
400 Maryland Ave. SW
Washington, DC 20202

1 Dr. Joseph L. Young, Director
Memory & Cognitive Processes
National Science Foundation
Washington, DC 20550

Private Sector

1 Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

1 Dr. Patricia Baggett
Department of Psychology
University of Colorado
Boulder, CO 80309

1 Mr. Avron Barr
Department of Computer Science
Stanford University
Stanford, CA 94305

1 Dr. Menucha Birenbaum
School of Education
Tel Aviv University
Tel Aviv, Ramat Aviv 69978
Israel

1 Dr. John Black
Yale University
Box 11A, Yale Station
New Haven, CT 06520

1 Dr. John S. Brown
XEROX Palo Alto Research Center
3333 Coyote Road
Palo Alto, CA 94304

1 Dr. Glenn Bryan
6208 Poe Road
Bethesda, MD 20817

1 Dr. Jaime Carbonell
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

1 Dr. Pat Carpenter
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

1 Dr. William Chase
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

1 Dr. Micheline Chi
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 Dr. Michael Cole
University of California
at San Diego
Laboratory of Comparative
Human Cognition - D003A
La Jolla, CA 92093

1 Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

1 Dr. Lynn A. Cooper
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 ERIC Facility-Acquisitions
4833 Rugby Avenue
Bethesda, MD 20014

1 Professor Reuven Feuerstein
HWCRI Rabov Karmon 6
Bet Hakerem
Jerusalem
Israel

1 Mr. Wallace Feurzeig
Department of Educational Technology
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02238

1 Dr. Dexter Fletcher
WICAT Research Institute
1875 S. State St.
Orem, UT 22333

1 Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

1 Dr. Don Gentner
Center for Human Information Processing
University of California, San Diego
La Jolla, CA 92093

1 Dr. Dedre Gentner
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02138

1 Dr. Robert Glaser
Learning Research & Development Center
University of Pittsburgh
3939 O'Hara Street
PITTSBURGH, PA 15260

1 Dr. Marvin D. Glock
217 Stone Hall
Cornell University
Ithaca, NY 14853

1 Dr. Joseph Goguen
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

1 Dr. Daniel Gopher
Faculty of Industrial Engineering
& Management
TECHNION
Haifa 32000
ISRAEL

1 DR. JAMES G. GREENO
LRDC
UNIVERSITY OF PITTSBURGH
3939 O'HARA STREET
PITTSBURGH, PA 15213

1 Dr. Barbara Hayes-Roth
Department of Computer Science
Stanford University
Stanford, CA 95305

1 Dr. Earl Hunt
Dept. of Psychology
University of Washington
Seattle, WA 98105

1 Dr. Marcel Just
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

BEST COPY AVAILABLE

1 Dr. David Kieras
Department of Psychology
University of Arizona
Tucson, AZ 85721

1 Dr. Walter Kintsch
Department of Psychology
University of Colorado
Boulder, CO 80302

1 Dr. Stephen Kosslyn
1236 William James Hall
33 Kirkland St.
Cambridge, MA 02138

1 Dr. Pat Langley
The Robotics Institute
Carnegie-Mellon University
Pittsburgh, PA 15213

1 Dr. Jill Larkin
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

1 Dr. Alan Lesgold
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

1 Dr. Jim Levin
University of California
at San Diego
Laboratory for Comparative
Human Cognition - D003A
La Jolla, CA 92093

1 Dr. Michael Levine
Department of Educational Psychology
210 Education Bldg.
University of Illinois
Champaign, IL 61801

1 Dr. Marcia C. Linn
Lawrence Hall of Science
University of California
Berkeley, CA 94720

1 Dr. Don Lyon
AFHRL/OT (UDRI)
Williams AFB, AZ 85225

1 Dr. Jay McClelland
Department of Psychology
MIT
Cambridge, MA 02139.

1 Dr. James R. Miller
Computer*Thought Corporation
1721 West Plano Highway
Plano, TX 75075

1 Dr. Mark Miller
Computer*Thought Corporation
1721 West Plano Parkway
Plano, TX 75075

1 Dr. Tom Moran
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304

1 Dr. Allen Munro
Behavioral Technology Laboratories
1845 Elena Ave., Fourth Floor
Redondo Beach, CA 90277

1 Dr. Donald A Norman
Cognitive Science, C-015
Univ. of California, San Diego
La Jolla, CA 92093

1 Dr. Jesse Orlansky
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA 22311

1 Prof. Seymour Papert
20C-109
Massachusetts Institute of Technology
Cambridge, MA 02139

1 Dr. Nancy Pennington
University of Chicago
Graduate School of Business
1101 E. 58th St.
Chicago, IL 60637

1 DR. PETER POLSON
DEPT. OF PSYCHOLOGY
UNIVERSITY OF COLORADO
BOULDER, CO 80309

1 Dr. Fred Reif
Physics Department
University of California
Berkeley, CA 94720

1 Dr. Lauren Resnick
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 1521

1 Dr. Jeff Richardson
Denver Research Institute
University of Denver
Denver, CO 80208

1 Mary S. Riley
Program in Cognitive Science
Center for Human Information Processing
University of California, San Diego
La Jolla, CA 92093

1 Dr. Andrew M. Rose
American Institutes for Research
1055 Thomas Jefferson St. NW
Washington, DC 20007

1 Dr. Ernst Z. Rothkopf
Bell Laboratories
Murray Hill, NJ 07974

1 Dr. William B. Rouse
Georgia Institute of Technology
School of Industrial & Systems
Engineering
Atlanta, GA 30332

1 Dr. David Rumelhart
Center for Human Information Processing
Univ. of California, San Diego
La Jolla, CA 92093

1 Dr. Michael J. Samet
Perceptronics, Inc
6271 Varial Avenue
Woodland Hills, CA 91364

1 Dr. Roger Schank
Yale University
Department of Computer Science
P.O. Box 2158
New Haven, CT 06520

1 Dr. Walter Schneider
Psychology Department
603 E. Daniel
Champaign, IL 61820

1 Dr. Alan Schoenfeld
Mathematics and Education
The University of Rochester
Rochester, NY 14627

1 Mr. Colin Sheppard
Applied Psychology Unit
Admiralty Marine Technology Est.
Teddington, Middlesex
United Kingdom

1 Dr. H. Wallace Sinaiko
Program Director
Manpower Research and Advisory Services
Smithsonian Institution
801 North Pitt Street
Alexandria, VA 22314

1 Dr. Edward E. Smith
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

1 Dr. Richard Snow
School of Education
Stanford University
Stanford, CA 94305

1 Dr. Elliott Soloway
Yale University
Department of Computer Science
P.O. Box 2158
New Haven, CT 06520

1 Dr. Kathryn T. Spoehr
Psychology Department
Brown University
Providence, RI 02912

1 Dr. Robert Sternberg
Dept. of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520

1 Dr. Albert Stevens
Bolt Beranek & Newman, Inc.
10 Moulton St.
Cambridge, MA 02238

1 David E. Stone, Ph.D.
Hazeltine Corporation
7680 Old Springhouse Road
McLean, VA 22102

1 Dr. Kikumi Tatsuoka
Computer Based Education Research Lab
252 Engineering Research Laboratory
Urbana, IL 61801

1 Dr. Perry W. Thorndyke
Perceptronics, Inc.
545 Middlefield Road, Suite 140
Menlo Park, CA 94025

1 Dr. Douglas Towne
Univ. of So. California
Behavioral Technology Labs
1845 S. Elena Ave.
Redondo Beach, CA 90277

1 Dr. Kurt Van Lehn
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304

1 Dr. Keith T. Wescourt
Perceptronics, Inc.
545 Middlefield Road, Suite 140
Menlo Park, CA 94025

1 William B. Whitten
Bell Laboratories
2D-610
Holmdel, NJ 07733

BEST COPY AVAILABLE